

## Patent

FOR

## Digital Content Distribution

Vaughn Iverson  
Todd Schwartz

[illegible]

Prepared by

Steven D. Yates  
Reg. No. 42,242  
(503) 264-6589

Express Mail mailing label number: EL034437237US

## Digital Content Distribution

### 5 **Field of the Invention**

The invention generally relates to distribution of digital content, and more particularly to creating, a manifest having structure, meta-data describing digital content, and references to digital resources, where the manifest facilitates finding, obtaining, organizing, collecting, and sharing the digital content.

10

### **Background**

Computing devices, such as a personal computer, typically employ a folder/document metaphor for storing and organizing content. While such a metaphor is well suited to an office environment, where content is typically in the form of documents, it is not effective when applied to the use and management of arbitrary digital content. For example, a digital photo album may be defined to correspond to a physical photo album. Thus, certain photos and/or textures may be identified as corresponding to front and back covers for the photo album, and other photos identified as being internal to the album. All content within the album would be stored in a folder or hierarchy of folders.

20 Unfortunately, such a hierarchy lacks cohesion or inherent structure found in a physical photo album. Thus, while a real photo album can be easily organized, shared, etc., there is no standard method for doing so with the hierarchy for the digital photo album.

One solution has been to apply an encapsulating structure (e.g. zip or tar file) to the hierarchy and thus create an archive containing all of the digital content. This archive can then be exchanged, sold, etc. A significant drawback to such an archive is

that it is monolithic, e.g., it contains all data within the album. This reduces the transferability of the album, and restricts selective transfer of only some album contents. In addition, an archive generally lacks data describing the context and contents of the photo album (e.g., meta-data), thus making it difficult to organize and search/find digital content, and to avoid having duplicative digital content. For example, once the archive it received, it needs to be re-created in a folder structure corresponding to the original hierarchy structure, or risk internal references to files being lost. This prevents or restricts moving, organizing, and re-organizing digital content.

In addition, assuming the user can maintain the proper structure of the content, also lacking to such an archive is an easy, robust, and widely-supported method for searching digital content. For example, for music data, without excessive redundancy, folder organization techniques cannot help if you want to locate a song played based on multiple criteria, such as mood and music type, since music would generally fall into multiple categories and thus require entry in multiple folders.

Similarly, one cannot select a video starring a certain actor and start viewing at a particular scene, hum a bar of a tune and have the system find it for you, or find all photos taken at an event last year. These tasks cannot be performed because mere hierarchical storage of digital content lacks associative or context data, such as captions, dates, places, copyrights, artist, author, genre, etc. to allow identification of digital content based on such search criteria.

Assuming the digital content can be located, exchanging the digital content is difficult due to a lack of open and flexible standard to do so. Existing digital content packaging schemes are proprietary and tailored to specific content types, such as for

electronic books or music. Currently there is no technology for uniformly packaging, in electronic form, a wide variety of digital content in a way that qualifies the content as what the entertainment industry calls a "title."

## 5 **Brief Description Of The Drawings**

The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

FIG. 1 illustrates a generalized network environment according to one embodiment of the invention.

10 FIG. 2 illustrates one embodiment of a manifest for representing digital content corresponding to a music album.

FIG. 3 illustrates, according to one embodiment of the invention, portions of an XML skeleton for implementing the FIG. 2 manifest.

15 FIG. 4 is a flowchart illustrating, according to one embodiment of the invention, creating digital content for use by a receiver.

FIG. 5 is a flowchart illustrating, according to one embodiment of the invention, receiving a manifest or manifests from an agent, such as in response to a search query.

FIG. 6 illustrates a suitable computing environment in which certain aspects of the invention may be implemented.

20

## **Detailed Description**

To address limitations in management and distribution of digital content, a system may be used for uniform electronic packaging of digital content independent of

content type, along with automated distribution of digital content to desired recipients.

The phrase “digital content” refers to digital representations of created works, such as multimedia items including photo albums, music albums, novels, magazines, and the like, as well as entirely new forms of content that may emerge in the electronic realm.

- 5 The phrase “electronic package” means an instance of digital content that exists in some computer accessible medium, such as on a wire, in computer memory, on a storage device, etc.

The illustrated embodiments of the invention may represent, or package, abstract created works in a concrete electronic digital form, as well as provide an electronic  
10 analog to traditional physical goods. For example, a CD-ROM or vinyl album is simply a physical representation of the abstract concept “music album”; a book is simply a physical representation of the abstract concept “novel.” But, it will be appreciated that these exemplary album and novel concepts have been largely defined with respect to the limitations of their physical representations. And, these abstract forms can be  
15 faithfully represented in an electronic form. However, it should be appreciated that since a non-physical electronic realm lacks limitations of physical objects, and therefore abstract forms of created works can morph and evolve in ways unpredictable, and perhaps not physically possible.

Thus, an electronic package may be structured to correspond to, or mimic, a  
20 physical good. For example, an electronic package corresponding to a music compact disc (CD) can be defined as having packaging, cover illustrations, organization, etc., corresponding to a physical CD. However, unlike the physical good, in an electronic realm, the electronic package may also include meta-data describing the parts, and/or

context of the packaged digital content, rules controlling distribution of the digital content, as well as references to content related to but not part of the physical good.

In one embodiment, an electronic package of digital content is represented by a “manifest” (see FIG. 2) that appears as a single unit, or item, to a user/receiver of the manifest. The manifest comprises “component” definitions corresponding to components of the digital content, each component definition specifying a reference to a digital resource.

FIG. 1 illustrates a general computing environment according to one embodiment of the invention. As illustrated, a manifest for an electronic package is created or transmitted by an originator **100-104**, and received by a receiver **106, 108**. It will be appreciated that the terms “receiver” and “originator” are arbitrary, in that one may also perform the operations of the other. Agents **110, 112** facilitate transfer of manifests from originators to receivers. Originators, receivers, and agents are in communication with each other over a network **114** such as an intranet, the Internet, or other network. There may be other networks, not illustrated, providing alternate or direct connections between the agents, originators, and/or receivers.

Agents include, for example, search engines used by a receiver to locate digital content of interest to the receiver, order processing systems for responding to and fulfilling purchase requests, and systems that customize content en route for specific locales. Agents can incorporate arbitrary “intelligence” that facilitate buying, selling, trading, bartering, lending, etc. of digital content. In one embodiment, originators and receivers may also perform the functions of an agent.

Thus, receiver 1 **106** may contact Agent 1 **110**, and indicate interest in obtaining a certain classical music CD as digital content. If Agent 2 cannot directly provide the requested digital content, Agent 1 contacts Agent 2 **112** over the network **114**, and request the desired digital content. Assuming Agent 2 can fulfill the request, Agent 2  
5 sends Agent 1 a manifest corresponding to the requested digital content.

It will be appreciated that typical content distribution services, such as Universal Music Group's bluematter and BMG (both online music providers), and Flycode (a peer-to-peer short film distribution network), already allow one to obtain digital content over a network. (Please note that all marks herein are the property of their respective owners.)

10 However, a significant limitation of such services is that received content is a monolithic encapsulation, it contains all parts of the digital content whether one wishes to receive it or not. Thus, for a music album, the receiver would receive all data for all songs within the album, which may potentially include multiple encodings of each song, etc. This contrasts distribution by way of a manifest, which may only contain references to digital  
15 resources, allowing advance determination of which parts, if any, of the content to obtain before expending resources to receive it.

Another limitation of prior art techniques is that the content is rigidly structured. For example, bluematter and BMG can provide a "song," or an "album," but these are content forms that cannot be easily extended to include other types of content, such as  
20 video clips, discount coupons, etc. Such content is also inflexible in the types of metadata that can be included. For example, fixed file formats (e.g., MP3, WMA, etc.) typically used in such services have specific notions of what constitutes appropriate metadata for a music album or an individual song. This results in limitations in what

data may be represented, and how, thus limiting general applicability of the file format. Such systems are tightly bound with specific encoding formats, and are difficult to upgrade when technological improvements are made.

5           FIG. 2 illustrates one embodiment of a manifest **200** for representing digital content corresponding to an abstract music album. It will be appreciated, however, that a manifest may also be used to electronically represent particular physical implementations of the abstract idea, e.g., to represent compact discs, vinyl albums, etc. It is assumed the manifest is created with an appropriate manifest generation /  
10           authoring tool. The manifest binds together various types of resources in their original authored formats. In one embodiment, the manifest only references data that can be created with well known or standard data authoring tools, such as MP3 generators, video editing software, etc.

15           As illustrated, the manifest **200** itself corresponds to a top, or outer, level. Within the illustrated manifest are six item declarations corresponding to two songs **202**, **216**, two videos **230**, **236**, cover art graphics **242** and meta-data **244**. The meta-data **244** may contain distribution rules, requirements, or policies (hereafter "rules or policies") regarding handling of the manifest, or the digital content described by the manifest.

20           Items (and sub-items) can have sub-items. For example, as illustrated, song 1 **202** contains references **204**, **206** to different encodings of song 1, meta-data **208** describing song 1, a reference **210** to lyrics for song 1, a reference **212** to a musical instrument digital interface (MIDI) encoding of song 1. The references **204**, **206** to the different encodings allow providing different encodings of a song, such as a high bit rate

version and a low bit rate version, allowing a receiver **106, 108** (FIG. 1) to elect a version based on certain criteria. For example, a receiver may elect a version based on communication data rates available to the receiver. The meta-data **208** for song 1 **202** may include song data such as playing time, composer, recording date, copyright information, and the like.

The meta-data for song 1 **202** may also contain pricing information for each song encoding, e.g., the high bit rate encoding **206** for song 1 may cost more than the low bit rate **204** encoding. The meta-data may also include other data such as content ratings, intended audience ratings, e.g., an indicator the song is intended for kids, teenagers, adults, etc., song genre, geographic distribution rules, e.g., different geographic regions receive different song versions, parental controls, etc. The manifest primarily contains references to digital resources, rather than directly embedding the resources. The meta-data may be used to determine how to handle digital resources referenced by the manifest (see FIG. 5). The meta-data can also be inspected to prevent duplicative retrieval of digital resources already be present in a cache or local storage, e.g., saved to disk or a collection. Duplicative retrievals unnecessarily burden networks and increase purchase costs and delivery times.

The illustrated manifest **200** also includes videos sections **230, 236** having references **232, 238** to video clips of concerts, music videos, or other video data related to the physical good described by the manifest. As with the songs, **202, 216**, the video sections **230, 236** may have associated meta-data **234, 240**, such as title, running time, content ratings, and other such data as discussed above for song data. It will be appreciated that video data is not ordinarily part of a music CD distribution. However,

since the manifest contains references to digital resources, rather than the resources themselves, it is easy to extend the manifest to include varied data typically associated with the music CD, even though such an inclusion would be impractical or cost ineffective to distribute along with the physical CD, or obtain along with a traditional  
5 monolithic download.

As illustrated, the manifest may also include a reference **242** to a graphics file, such as for a graphics encoding of the cover art for the music CD. Such graphics data facilitates presenting, by way of a user interface, a visual representation, such a picture of the cover of an album, or other representative view, of the created work being  
10 represented by the manifest.

It is assumed that an appropriate user interface, such as a text-based or graphical-based interface is used to identify digital resources to include in the manifest, and to create the manifest accordingly. In one embodiment, the same user interface is used to process received manifests. In one embodiment, the graphical user interface is  
15 based on an Internet browser. In one embodiment, the manifest includes security provisions to detect unauthorized tampering with the manifest. In one embodiment, the manifest is cryptographically signed. In another embodiment, some or all of the manifest is encrypted with a private and/or known public cryptographic systems. See, e.g., *Applied Cryptography : Protocols, Algorithms, and Source Code in C* by Bruce  
20 Schneier.

FIG. 3 illustrates, according to one embodiment of the invention, portions of an XML skeleton **300** for implementing the FIG. 2 manifest **200**. The eXtensible Markup

Language (XML) is used due to its widespread industry support and structures for organizing data. XML allows creation of custom tags, which enables defining, validating, and interpreting transmitted manifests. It will be appreciated, however, that other languages, such as tag based languages, or scripting languages, may also be used. In addition, in one embodiment, a binary data standard (e.g., cross-platform readable) may be used to encode a manifest. In general, any data format that allows data sets to be named, and hierarchically structured, can be used.

A first tag **302** indicates the beginning of a music package definition, or container, for the digital content defined by the manifest. The package has multiple defined items.

A first item **304** contains information about the "digital item" (in this case, an electronic version of a music album) described by the manifest, and configuration options regarding handling the contents of the package, e.g., how to prompt a receiver regarding processing songs and other digital content referenced by the manifest.

Each defined item, and the manifest itself, may have arbitrary meta-data to allow associating real-world context with an item. Meta-data is introduced with descriptor tags **320, 322** having definitions **324**. A descriptor tag NAME attribute indicates the meta-data attribute being defined, and the VALUE attribute supplies an appropriate value.

The first item **304** includes a choice tag **306** identified with an ID **308** to allow a receiver of the manifest to choose how to receive songs referenced by the manifest. A selection tag **310**, e.g., "pick songs," provides an option for selecting songs individually. It will be appreciated that other selection options, or no options, may be provided.

A subsequent choice tag **312** indicates available songs, e.g., SONG1, SONG2, from which to choose. Note that the choice tag **312** has a condition tag **314** requiring

the above "pick songs" selection tag **310**. This condition must be satisfied before this choice **312** is acted upon, and it operates in the context of the previous choice tag **306**. In such fashion, manifests may be structured to provide choices based on previous selections. Subsequent selection tags **316** define songs that may be individually  
5 selected.

A following tag **318** indicates a start of a manifest section for selecting a desired bit rate desired for selected songs. As above with song selection, bit rate selection options (not illustrated) may be included to operate with respect to the method employed for song selection. For example, if an election is made to individually pick  
10 songs, then one may be prompted for each song to select a desired bit-rate.

Other choice tags may also be defined to control delivery of the digital content to another party, e.g., to include an address book to facilitate sale or transfer of digital content to another receiver. Tags may also be defined to prompt a receiver to identify a computing platform receiving the manifest, e.g., desktop computer, portable, hand-held  
15 device, personal digital assistant (PDA), as each such device may have different format requirements, storage capacities and interface capabilities.

Tags may also be defined to indicate the operating system in use by the receiving device, and to control whether to obtain extra content, such as related videos **230, 236**, cover art and related package graphics, lyrics, press reviews, etc. Tags may  
20 also be defined to indicate data format preferences, e.g., graphics images should be converted (if necessary) into Joint Photographic Experts Group (JPEG) format images, or videos should be in a particular format.

Following various choice tags, are descriptor tags **322** that are at the same hierarchical level as the first item **304** definition, e.g., a music compact disc, and therefore define **324** typical characteristics for the disc. Exemplary illustrated characteristics include the type of music disc defined by the item **304** (an album), an audience rating, album title, creator, contributor, album publisher, and album rights (e.g., copyright information). As discussed above, any of these attributes can be used to control or facilitate organization, content filtering, and re-distribution of digital content.

This description assumes the receiver is a person using an application program or device configured to receive and/or process the manifest, and display choices accordingly. However, note that the receiver may be another application program or device, such as one operating to automatically apply local policies or rules to received manifests, before a manifest is received by a person. In particular, there may be local policies, such as parental controls, and these policies may be used to cause a rewriting of the manifest, e.g., per XML rewriting rules, so that the manifest excludes references to undesirable content.

It will be appreciated that the exemplary FIG. 3 skeleton may be constructed with content and features to free a user from tedious download processes, and from the need to know or understand file formats or plug-ins (e.g., application program enhancements designed to process a particular data format) when receiving content from network originators **100-104** (FIG. 1) . Through use of agents **110, 112**, or direct network connections, the user can place an order for digital content, and then continue with other tasks. The order is fulfilled in the background, without user intervention. Delivery can be robust such that if a network connection is broken, a data transfer

continues automatically when the network connection is reestablished. When content has been received, and its manifest accepted, modified or rejected according to user and local-system rules and/or policies, the user can be notified accordingly.

5           FIG. 4 is a flowchart illustrating, according to one embodiment of the invention, creating an electronic package of digital content for use by a receiver. As illustrated, a first operation is to package digital content **400** for distribution to receivers; as discussed above, in one embodiment, a manifest is used to package the content. A packaging program is used to identify content to be packaged and generate a  
10 corresponding manifest. Individual digital resources can be identified by way of traditional file selection dialogs, or with Object Linking and Embedding (OLE) type drag and drop operations, or through other selection techniques.

When created, the manifest contains references to identified resources, and if desired, a virtual structure corresponding to a physical good, e.g., a music CD. Thus, a  
15 receiving program for a music CD manifest would present a user interface displaying the virtual representation of the physical good. In the CD example, this might include rendering a CD jewel box with appropriate cover art and song listings as defined within the manifest.

A created manifest is sent **402** to an agent for distribution, which in turn stores  
20 **404** the manifest in a public storage. Note that the storage need not be publicly accessible, however public access is assumed so that search engines, agents, and other location applications are able to mine and/or index stored manifests. It is assumed that the agent stores the manifest in a database. The digital resources

identified by the manifest are stored **406** as well. The digital resources may be stored in public or private storage locations depending on intellectual property rights for the digital resources.

Once the manifest has been created and stored, an agent can receive **408** a search query, e.g., purchase request, database look-up request, etc. The agent searches **410** through the database for manifests corresponding to the query.

Assuming manifests are XML based, it will be appreciated that different query languages may be used depending on manifest formatting and database technology used to store the manifests. For example, the agent may use XML Query Language (XQL) queries, Stanford University's Lore query language, or other query languages.

Manifests matching the received **408** query are returned **412** to a requestor. This may result in a request to purchase digital content identified by a manifest. Towards this end, the agent would confirm payment for digital content as required by a manifest, and make digital content available to a purchaser. Various known techniques may be employed to provide purchased digital content in a secure fashion so as to reduce possibility of illicit appropriation of the digital content, e.g., temporary names and locations for the digital resources included in the content, creation of secure communications channels or tunnels between the agent and purchases, or the like.

FIG. 5 is a flowchart illustrating, according to one embodiment of the invention, receiving **500** a manifest or manifests from an agent, such as in response to a search query. It is assumed the manifest is received by an application program, operating system extension, or the like, before a user / person is granted access to the manifest

or allowed to take action on (e.g., purchase) the digital content referenced by the manifest.

When the manifest is received, a receiver **106, 108** (FIG. 1) applies **502, 506** rules or policies to the meta-data **208, 222, 234, 240, 244** (FIG. 2) to determine how to process the digital content, such as whether to reject some or all of the manifest. For example, a receiver may have local rules controlling whether to include supplemental content such as the videos **230, 236**, (FIG. 2) MIDI encodings **212, 226** (FIG. 2), or other data not traditionally part of an original physical good, e.g., the music CD. A user preference may be set to cause discarding of any such supplemental manifest portions.

In one embodiment, there may be different rules or policies **506** for each user of the receiving device, e.g., the receiving device may be a multiple-user computer, in addition to global rules or policies that are applied **502** to all users of the computing device. For example, parental controls may be used to globally restrict the type of content that any user may receive, or content known to be incompatible with the receiver may be automatically edited out of a manifest. In one embodiment, a received manifest is rewritten **504**, if necessary, to comply with the global rules or policies. Note that a remote management system can be used to set or change rules or policies.

After checking the manifest against global rules or policies, user rules or policies are then applied **506** to the manifest, and the manifest again rewritten **508**, if necessary, to comply with the user rules or policies. In one embodiment, the process of “rewriting” a manifest is done in a rigorous manner. The rules controlling the rewriting are expressed in declarative form within the manifest itself within <choice> and <selection> tags, and their corresponding <condition> tags. This makes the rewriting process

robust, and ensures the resulting rewritten manifest is valid. Once a received manifest has cleared both global and user rules or policies, the resultant manifest is executed or interpreted 510 by the receiver. As discussed above with respect to FIG. 3, such processing may cause a user interface to be displayed to allow a user selection among 5 different choices provided by the manifest for buying or obtaining 512 digital content accordingly.

In one embodiment, manifests are encoded according to a rules based grammar comprising the following exemplary rules:

- |     |           |   |
|-----|-----------|---|
| (1) | container | ::= container <sup>*</sup> item <sup>*</sup> meta-data <sup>*</sup>                                   |
| (2) | item      | ::= (item   component) <sup>+</sup> choice <sup>*</sup> meta-data <sup>*</sup> condition <sup>*</sup> |
| (3) | component | ::= resource meta-data <sup>*</sup> anchor <sup>*</sup> condition <sup>*</sup>                        |
| (4) | meta-data | ::= meta-data <sup>*</sup> (component   statement) condition <sup>*</sup>                             |
| (5) | anchor    | ::= reference meta-data <sup>*</sup> condition <sup>*</sup>   |
| (6) | choice    | ::= selection <sup>+</sup> meta-data <sup>*</sup> condition <sup>*</sup>                              |
| (7) | selection | ::= predicate meta-data <sup>*</sup> condition <sup>*</sup>   |
| (8) | condition | ::= predicate <sup>+</sup>  |

Where: a container may be a hierarchical structure allowing items to be grouped to form logical packages that may be transported, exchanged, sold, stored, etc. An item is a grouping of sub-items and/or components associated with relevant meta-data about the item. Items may contain choices for customizing and/or configuring the item, and may be conditioned on predicates asserted by selections defined in the choices. A resource is a collection of data, such as an audio clip, image file, text document, or other data. A component is an association between a resource and all of its relevant meta-data; component meta-data will typically contain control or structural information about a resource, e.g., bit rate, media type, character set, encryption information, etc. A reference designates a portion of a resource, such as a specific location or range within the resource. A "statement" comprises a value and an associated domain for that

value. And, a predicate is a non-substitutable token, and in one embodiment, has a true, false, or undecided value.

FIG. 6 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which certain aspects of the illustrated invention may be implemented.

An exemplary system for implementing the invention includes a computing device **600** having system bus **602** for coupling various computing device components. Typically, attached to the bus are non-programmable and programmable processors **604**, a memory **606** (e.g., RAM, ROM), storage devices **608**, a video interface **610**, and input/output interface ports **612**. Storage devices include hard-drives, floppy-disks, optical storage, magnetic cassettes, tapes, flash memory cards, memory sticks, digital video disks, and the like.

The invention may be described by reference to different high-level program modules and/or low-level hardware contexts. Those skilled in the art will realize that program modules can be interchanged with low-level hardware instructions. Program modules include procedures, functions, programs, components, data structures, and the like, for performing particular tasks or implementing particular abstract data types. Modules may be incorporated into single and multi-processor computing devices, Personal Digital Assistants (PDAs), cellular telephones, and the like. Thus, the storage systems and associated media can store data and executable instructions for the computing device.

As discussed above with respect to FIG. 1, the computing device is expected to operate in a networked environment using logical connections to one or more remote computing devices **614**, **616** through a network interface **618**, modem **620**, or other communication pathway. Computing devices may be interconnected by way of a network **622** such as an intranet, the Internet, or other network. Modules may be implemented in a single computing device, or processed in a distributed network environment, and stored in both local and remote memory. Thus, for example, with respect to the illustrated embodiments, assuming computing device **600** is a computing device operated by a user searching for music in a particular genre, then remote devices **614**, **616** may respectively be first and second agents that receive and respond to search issued by computing device **600**.

It will be appreciated that remote computing devices **614**, **616** may be configured like computing device **600**, and therefore include many or all of the elements discussed for computing device. It should also be appreciated that computing devices **600**, **614**, **616** may be embodied in a single device, or separate communicatively-coupled components, and may include or be embodied in routers, bridges, peer devices, web servers, and application programs utilizing network application protocols such as the HyperText Transfer Protocol (HTTP), File Transfer Protocol (FTP), and the like.

Having described and illustrated the principles of the invention with reference to illustrated embodiments, it will be recognized that the illustrated embodiments can be modified in arrangement and detail without departing from such principles.

And, even though the foregoing discussion has focused on particular embodiments, it is understood that other configurations are contemplated. In particular,

even though expressions such as "in one embodiment," "in another embodiment," or the like are used herein, these phrases are meant to generally reference embodiment possibilities, and are not intended to limit the invention to particular embodiment configurations. As used herein, these terms may reference the same or different  
5 embodiments, and unless implicitly or expressly indicated otherwise, embodiments are combinable into other embodiments. Consequently, in view of the wide variety of permutations to the above-described embodiments, the detailed description is intended to be illustrative only, and should not be taken as limiting the scope of the invention.

What is claimed as the invention, therefore, is all such modifications as may  
10 come in the scope and spirit of the following claims and equivalents thereto.